

HP OpenView Storage Mirroring

Evaluation Guide



© Copyright 2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Microsoft®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

HP OpenView Storage Mirroring Evaluation Guide

Sixth Edition (June 2005)

Part Number: T2558-90008

NSI and Double-Take are registered trademarks of NSI Software, Inc. All other products are trademarks of their respective companies.

(c) 1996-2005 NSI Software, Inc.

Table of Contents

Chapter 1: Introduction	1-1
Evaluation guide overview	1-1
Storage Mirroring Operations	1-2
Mirroring	1-2
Replication	1-3
Failure monitoring/failover	1-4
Restoration	1-5
Chapter 2: Resources	2-1
Chapter 3: Storage Mirroring Setup	3-1
System requirements	3-1
Installing Storage Mirroring	3-2
Chapter 4: Storage Mirroring Evaluation	4-1
Creating initial test data	4-1
Creating and connecting a replication set	4-2
Monitoring the activity and completion of the initial mirror	4-7
Changing the test data to cause replication	4-8
Verifying the changes on the target server	4-10
Visual verification	4-10
Third-party verification	4-10
Storage Mirroring verification	4-11
Configuring failover	4-12
Monitoring failover	4-14
Initiating a failure	4-15
Performing failback	4-16
Restoring data	4-18
Chapter 5: Conclusion	5-1

I ► Introduction

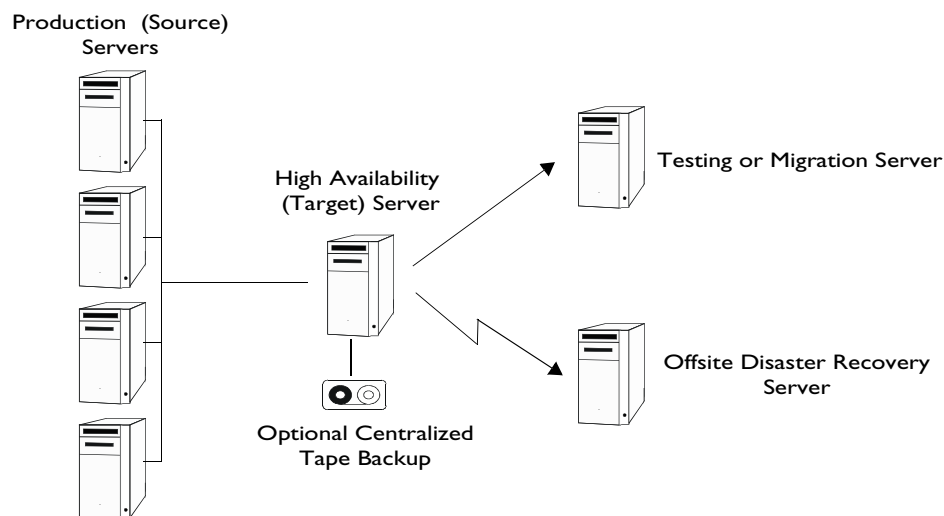
Thank you for your interest in evaluating HP OpenView Storage Mirroring®. The purpose of this evaluation is to allow you to become familiar with the Storage Mirroring core replication and failover technology first hand. Storage Mirroring is a real-time data replication and failover software product. Storage Mirroring augments your existing data protection strategy by reducing downtime and data loss, and it provides these services with minimal impact on existing network and communication resources.

Storage Mirroring is a real-time data replication and failover software product. Storage Mirroring augments your existing data protection strategy by reducing downtime and data loss, and it provides these services with minimal impact on existing network and communication resources.

Storage Mirroring allows you to specify mission-critical data that must be protected and replicates, in real-time, that data from a production machine, known as the source, to a backup machine, known as the target. The target machine, on a local network or at a remote site, stores the copy of the critical data from the source. Storage Mirroring monitors any changes to the critical data and sends the changes to the target machine. By replicating only the file changes rather than copying an entire file, Storage Mirroring allows you to more efficiently use resources.

Storage Mirroring lets you implement various data protection solutions including:

- ◆ Local high availability services
- ◆ Offsite disaster recovery services
- ◆ Enhanced centralized backup using third-party backup systems



Sample Configuration

Evaluation guide overview

In this guide, you will find five sections:

1. **Introduction**—A brief overview of Storage Mirroring
2. **Resources**—Resources available during your evaluation
3. **Installation**—System requirements and step-by-step instructions for installing Storage Mirroring
4. **Operation**—Step-by-step instructions for basic Storage Mirroring operations
5. **Conclusion**—Final words on your Storage Mirroring evaluation

Storage Mirroring Operations

Storage Mirroring performs four basic types of operations:

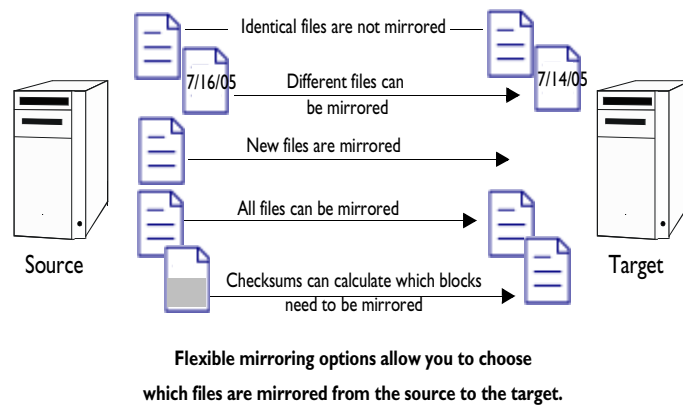
- ◆ **Mirroring**—The initial copy or subsequent resynchronization of selected data
- ◆ **Replication**—The on-going capture of byte-level file changes
- ◆ **Failure monitoring and failover**—The ability to monitor and stand-in for a machine, in the event of a failure
- ◆ **Restoration**—A mirror of selected data from the target back to the source

Each operation is briefly described in the following sections. For complete details, see the corresponding chapter.

Mirroring

Mirroring is the process of transmitting user-specified data from the source to the target so that an identical copy of data exists on the target. When Storage Mirroring initially performs mirroring, it copies all of the selected data including file attributes and permissions. Mirroring creates a foundation upon which Storage Mirroring can efficiently update the target machine by replicating only file changes.

If subsequent mirroring operations are necessary, Storage Mirroring can mirror specific files or blocks of changed data within files. By mirroring only files that have changed, network administrators can expedite the mirroring of data on the source and target machines.




Mirroring has a defined end point - when all of the selected files from the source have been transmitted to the target. When a mirror is complete, the target contains a copy of the source files at that point in time.

Replication

Replication is the real-time transmission of file changes. Unlike other related technologies, which are based on a disk driver or a specific application, the Storage Mirroring replication process operates at the file system level and is able to track file changes independently from the file's related application. In terms of network resources and time, replicating changes is a more efficient method of maintaining a real-time copy of data than copying an entire file that has changed.

After a source and target have been connected through Storage Mirroring, file system changes from the user-defined mission-critical data are tracked. Storage Mirroring immediately transmits these file changes to the target machine. This real-time replication keeps the data on the target machine up-to-date with the source machine and provides high availability and disaster recovery with minimal data loss.



Unlike mirroring which is complete when all of the files have been transmitted to the target, replication captures the changes as they are written to the source and sends them to the target continuously. Replication keeps the target up-to-date and synchronized with the source.


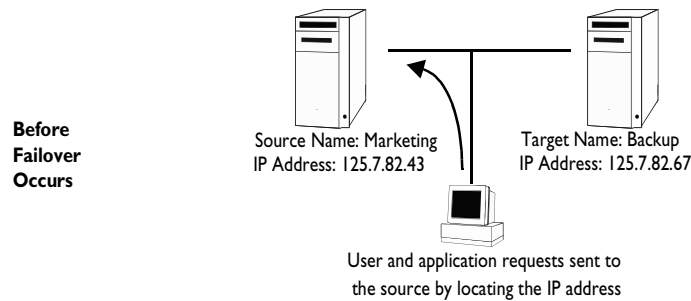
Failure monitoring/failover

Failover is a process in which a target machine stands in for a failed source machine. As a result, user and application requests that are directed to the failed source machine are routed to the target machine.

Storage Mirroring monitors the status of machines by tracking network requests and responses exchanged between source and target machines. When a monitored machine misses a user-defined number of requests, Storage Mirroring assumes that the machine has failed. Storage Mirroring then prompts the network administrator to initiate failover, or, if configured, it occurs automatically.

The failover target assumes the network identity of the failed source. When the target assumes the identity of the source, user and application requests destined for the source machine or its IP address(es) are routed to the target. After failover, user and application requests directed to the source machine are routed to the target machine.

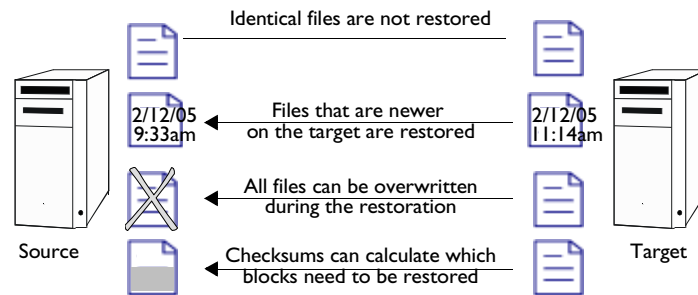
When partnered with the Storage Mirroring data replication capabilities, failover routes user and application requests with minimal disruption and little or no data loss. In some cases, failover may be used without data replication to ensure high availability on a machine that only provides processing services, such as a web server.



Restoration

Restoration provides an easy method for copying replicated data from the target back to its original location on the source. The process only requires you to select the source, target, and the appropriate replication set. There is no need to select files or to remember where the data came from on the source since that information is maintained by Storage Mirroring.

Restoration can be used if the source data is lost due to a disk crash or when the most up-to-date data exists on the target due to failover. At the time of a source machine failure, your Storage Mirroring target will contain the same data as your Storage Mirroring source. If you are using the Storage Mirroring failover capabilities, users can continue updating data on the target machine while the problems on the source are resolved. Because of the continued updates on the target, when the source machine is ready to come back online, the two machines will no longer contain the same data. Restoration is the process of copying the up-to-date data from the target back to the original source or new source location when bringing the source back online.



Flexible restoration options allow you to choose which files are restored from the target to the source.

Like mirroring, restoration has a defined end point - when all of the selected files from the target have been transmitted back to the source. When a restoration is complete, the source and target are again synchronized.

2 ► Resources

During the evaluation you have access to the same reference and online resources as licensed Storage Mirroring users.

- ◆ **Getting Started guide**—This reference manual is a quick reference guide for the installation and basic steps to get you going quickly with Storage Mirroring. It is available in Adobe® Acrobat® PDF format in the directory where you installed Storage Mirroring. The file is called `instlwin.pdf`.
- ◆ **User's Guide**—This reference manual is the complete reference guide with detailed steps for using Storage Mirroring. It is also available in Adobe Acrobat PDF format in the directory where you installed Storage Mirroring. The file is called `dtuser.pdf`.
- ◆ **Online help**—The Management Console and Failover Control Center online help requires at least a version 4.0 browser (version 4.5 is recommended). The online help can be accessed by pressing the F1 key, clicking the **Help** button on screens where it is available, or by selecting **Help, Help Topics**.

NOTE: Online Help is also available for the Text Client, but for the purpose of this evaluation, only the Management Console and Failover Control Center will be used.

If you need assistance during your evaluation, you can contact HP Customer Support. All basic configurations outlined in the Storage Mirroring *User's Guide* will be supported through Customer Support.

3 ► Storage Mirroring Setup

Your Storage Mirroring setup consists of two tasks: determining that your source and target meet the system requirements of this evaluation and performing the Storage Mirroring installation.

System requirements

Storage Mirroring is an exceptionally flexible product that can be used in a wide variety of network configurations. In a production environment, the multiple configurations can be used independently or in varying combinations. For the purpose of this evaluation, you should complete the evaluation procedures using two servers in a one-to-one, active/standby configuration. This scenario dedicates one target machine, having no production activity, to support one source machine. The source is the only machine actively replicating. This configuration is appropriate for disaster recovery, failover, and critical data backup.

To implement a one-on-one configuration, make sure that you have two servers that each meet the following system requirements:

- ◆ **Operating System and Licensing**—There are different Storage Mirroring licenses depending on the operating system you are using. Be sure you have the correct license for your operating system.

License	Valid Operating System
WNE WNE Replication ^a	<ul style="list-style-type: none">◆ Windows® Storage Server 2003 (Server or Advanced Server Edition) with SAK 3.0◆ Windows 2000 Powered OS (Server or Advanced Server Edition) with SAK 2.0 or 2.5
Server Edition	<ul style="list-style-type: none">◆ Windows Server 2003 Standard Edition◆ Windows Server 2003 Web Edition◆ Windows 2000 Server◆ Windows NT 4.0 Server◆ Windows NT 4.0 Terminal Server◆ Any valid operating system which can be used with a Storage Mirroring WNE license
Advanced Edition	<ul style="list-style-type: none">◆ Windows Server 2003 Enterprise Edition◆ Windows 2000 Advanced Server◆ Windows NT 4.0 Enterprise Edition◆ Any valid operating system which can be used with a Storage Mirroring Server Edition license
Datacenter Edition	<ul style="list-style-type: none">◆ Windows Server 2003 Datacenter Edition◆ Windows 2000 Datacenter Server◆ Any valid operating system which can be used with a Storage Mirroring Advanced Edition license

a. There are two WNE (Workgroup NAS Edition) licenses which run on the same operating systems. The WNE license contains full functionality while the WNE Replication license is limited to replication only. There is no failover functionality available with the WNE Replication license.

NOTE: Any version of Windows NT 4.0 must have Service Pack 4 or higher.

- ◆ **Storage Mirroring Version Interoperability**—Only particular Storage Mirroring versions (not editions) can interoperate. If both of your source and target are running the same Storage Mirroring version (for example, both are running version 4.4) there are no limitations.
- ◆ **Network**—TCP/IP LAN connection with one network adapter and one IP address

NOTE: Storage Mirroring servers can have multiple NICs and/or IP addresses. This limitation is only applied to simplify the evaluation process.

-
- ◆ **Name Resolution**—Your target server and client machine must be configured to use the same primary WINS server

NOTE: This limitation is also applied to simplify the evaluation process.

- ◆ **System Memory**—64 MB minimum

- ◆ **Disk Space**

Program Files	48 MB
Microsoft® Windows Installer Files	20 MB
Total	68 MB

NOTE: The program file can be installed to any volume while the Microsoft Windows Installer files are automatically installed to the operating system boot volume.

During the installation, you will be identifying how much disk space to use for Storage Mirroring queuing. This disk space is in addition to the total disk space above.

If the server is a target storing replicated data, it needs sufficient disk space to store replicated data from all connected sources. Be sure to allow additional space for growth, if needed.

- ◆ **E-mail Notification**—In order to enable Storage Mirroring e-mail notification, the server must have Internet Explorer 5.0 or later.

Additionally, for this evaluation, you will be using a third, client machine where you will be running the Storage Mirroring Management Console client and manipulating the server data. This machine can run either a Windows server operating system (Windows 2000 or Windows NT 4.0) or a workstation operating system (Windows ME or 98).

NOTE: To install Storage Mirroring you will need an activation code, which is included in your product packaging. If you downloaded the product from the web or do not have product packaging, you must call HP Customer Support.

Installing Storage Mirroring

Use these instructions to install Storage Mirroring on a machine for the first time. Complete installation details are available in the Storage Mirroring *Getting Started* guide.

1. Start the installation using one of the following two methods, depending on the source of your software:
 - ◆ **Storage Mirroring CD**—Insert the Storage Mirroring CD and the installation program starts automatically.
 - ◆ **Web download**—Run `setup.exe` from the location where you downloaded the Storage Mirroring files.
2. Click through the Welcome screen and then review and accept the Storage Mirroring license agreement to continue with the installation program.
3. Select the **Client and Server Components** installation option.
4. Select the **Complete** Storage Mirroring installation.
5. Enter your **User Name**, **Organization**, and **Activation Code**. The code is a 16-character, alpha-numeric activation code which applies the appropriate Storage Mirroring license to your installation.

NOTE: If you do not have an activation code, you must call HP Customer Support.

6. Select the remaining default selections and click **Install** to complete the installation.
7. After the installation, you must reboot the system prior to starting Storage Mirroring.

After the installation is complete, repeat these instructions on the second server. After both servers are complete, install Storage Mirroring on the client machine. On the client machine, install the **Client Components Only** in step 3 which skips the remaining choices.

4 ► Storage Mirroring Evaluation

Storage Mirroring configuration consists of creating a replication set, which defines the data on the source machine that is to be protected, and then connecting that replication set to a target machine. These two steps, creation and connection of the replication set, are the backbone of the Storage Mirroring processes. From there you can also use Storage Mirroring's failover, failback, and restoration processes for a high availability solution.

The evaluation consists of the following tasks:

1. *Creating initial test data*
2. *Creating and connecting a replication set*
3. *Monitoring the activity and completion of the initial mirror*
4. *Changing the test data to cause replication*
5. *Verifying the changes on the target server*
6. *Configuring failover*
7. *Monitoring failover*
8. *Initiating a failure*
9. *Performing failback*
10. *Restoring data*

NOTE: The last five sections of this evaluation cover failover, failback, and restore. These processes are used for high availability configurations. These sections are optional, although you are encouraged to step through the procedures.

Creating initial test data

1. Create a directory called `testdata` at the root of a volume on your source server.
2. Share the new directory by creating a share called `testshare`.
3. From your client machine, map `s:` to `testshare`.

```
net share testshare=d:\testdata
```

```
net use s: \\source\testshare
```

```
Microsoft Windows Advanced Server 2000  
(C) Copyright 1985-2000 Microsoft Corp.
```

```
D:\>md testdata
```

```
D:\>net share testshare=d:\testdata  
testshare was shared successfully.
```

```
D:\>net use s: \\saturn\testshare  
The command completed successfully.
```

4. Copy a significant amount (approximately 1 GB) of user data to your `s:` drive. This will be your test data that you will be replicating to your target.

NOTE: Make sure that your test data contains files that you can quickly view and edit, like a Microsoft Word or Notepad document.

Creating and connecting a replication set

The first time you start the Storage Mirroring Management Console (**Start, Programs, Storage Mirroring, Management Console**), the Welcome screen leads you to the Storage Mirroring Connection Wizard. The Connection Wizard will guide you through the essential steps necessary to establish a basic connection. In the Connection Wizard, you specify:

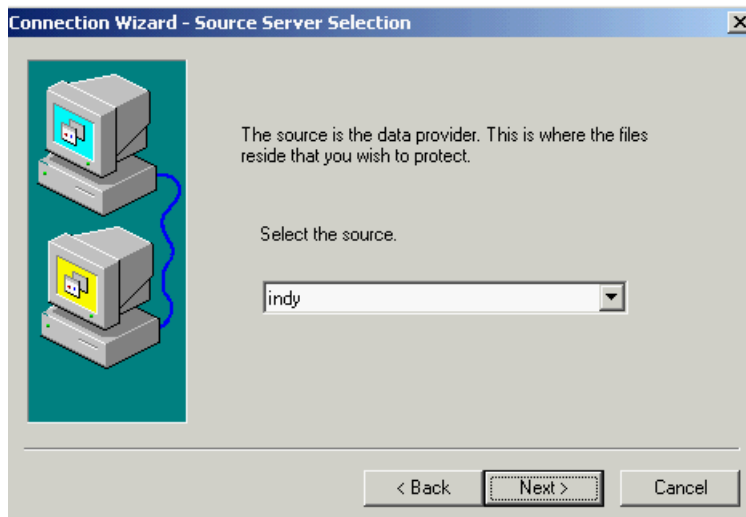
- ◆ **Source**—The machine containing the data to protect. This is your production server.
 - ◆ **Target**—The machine to receive the data from the source. This is your backup server.
 - ◆ **Replication Set**—The data to be protected.
 - ◆ **Target Location**—The location on the target machine to store the data.
 - ◆ **Advanced Options**—The Connection Wizard allows you to open the Connection Manager to set advanced options.
1. From your client machine, start the Management Console and select **Make your first connection** on the Welcome screen.

NOTE: If you have cancelled out of the Welcome screen, select **Tools, Connection Wizard**. (In order for the **Connection Wizard** menu option to be available, you must have a source or target highlighted in the left pane of the Management Console. You do not need to be logged on to the machine. If the Storage Mirroring Servers root is highlighted, the **Connection Wizard** option will not be available.)

You can also drag and drop a source machine on to a target machine on the left pane of the Management Console to start the Connection Wizard.

At any time while using the Connection Wizard, select **Back** to return to previous screens and review your selections.

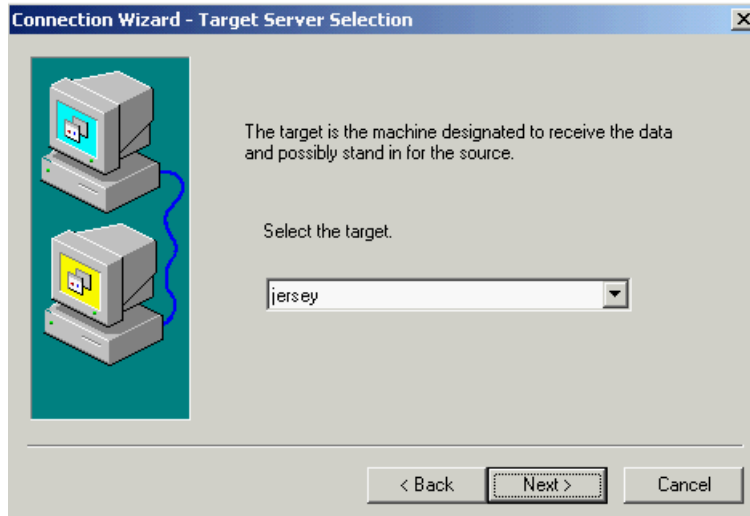
2. The Connection Wizard opens to the Welcome screen. Review this screen and click **Next** to continue.
3. Select your source machine from the drop down list. This is the machine where the files reside that you wish to protect. Click **Next** to continue.



NOTE: Storage Mirroring will automatically attempt to log on to the selected source using the identification of the user logged on to the local machine. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

The account specified must be a member of the machine's **Double-Take Admin** group. By default, members of the machine's local Administrators group are default members.

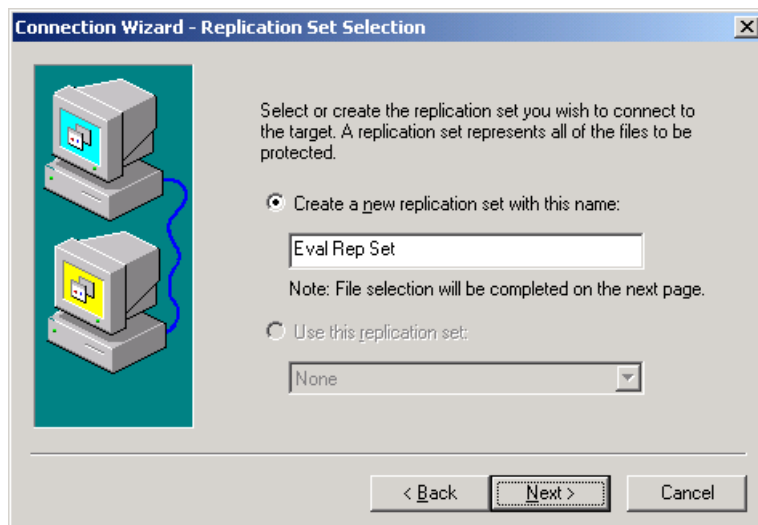
4. Select your target machine from the drop down list. This is your backup machine that will receive the data from the source. Click **Next** to continue.



NOTE: Storage Mirroring will automatically attempt to log on to the selected target using the identification of the user logged on to the local machine. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

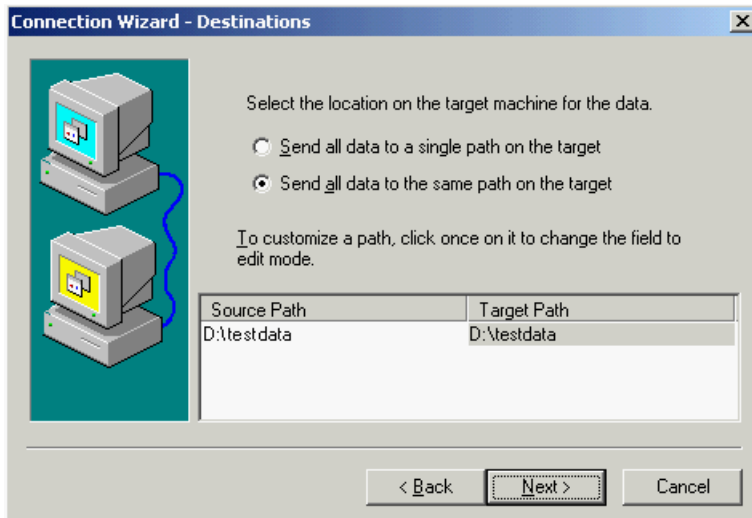
The account specified must be a member of the machine's **Double-Take Admin** group. By default, members of the machine's local Administrators group are default members.

5. Specify the name of the replication set, Eval Rep Set, that you want to create. Click **Next** to continue.



6. A tree display appears identifying the volumes and directories available on that source machine. Mark the `\testdata` directory that was created in the last section. This is the data that you will be protecting for this evaluation. Click **Next** to continue.

7. Select the option **Send all data to the same path on the target**. This option keeps the directory structure on the source and target identical. For example, d:\testdata is transmitted to d:\testdata.

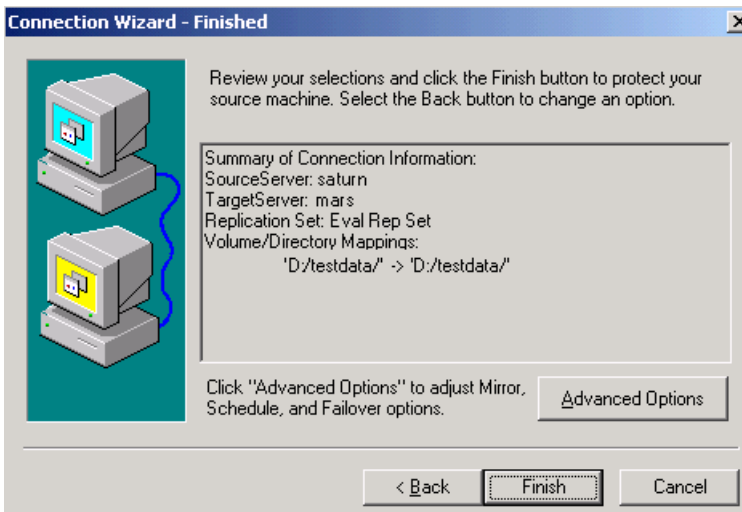


NOTE: This screen also allows you to **Send all data to a single path on the target** which sends the data to a default location which is defined as /source_name/rep_set_name/volume. For example, if you are replicating from your c: drive on a source machine called Saturn and your replication set is called Eval Rep Set, then your files will be located at /saturn/eval rep set/D on the target. This configuration is ideal if you are replicating multiple source servers to a single target.

You can also specify a custom location which allows you to specify exactly where you want the files located on the target.

For the purpose of this evaluation you should select the one-to-one, same path option.

8. Click **Next** to continue.
9. Review your selections on the summary screen.



10. After your Connection Wizard settings are correct, continue the evaluation by viewing the **Advanced Options**.


NOTE: You do not have to view the advanced options in order to establish a connection. You can click **Finish** and you will be prompted to save your replication set and when the Connection Wizard closes, the connection will be established, and mirroring and replication will begin immediately.

11. Click **Advanced Options** on the Connection Wizard summary screen, if you haven't already done so. The Connection Wizard will close and the Connection Manager will open. Notice that there are six tabs in the Connection Manager allowing you to control various advanced aspects of Storage Mirroring.

- ◆ **Servers**—The Servers tab controls the replication set and connection options that you established in the Connection Wizard.
- ◆ **Mirroring**—The Mirroring tab controls how the mirror process will be performed.
- ◆ **Orphans**—The Orphans tab configures whether orphan files, which are files that are present in the target location but not in the source replication set, will be deleted or moved.
- ◆ **Transmit**—The Transmit tab allows you to specify transmission limiting criteria including bandwidth throttling, start and stop criteria, and a transmission window.
- ◆ **Verify**—The Verify tab allows you to specify when a verification will be performed and how often.
- ◆ **Failover**—The Failover tab allows you to configure failure monitoring.
- ◆ **Compression**—The Compression tab allows you to set data compression levels for mirroring and replication.

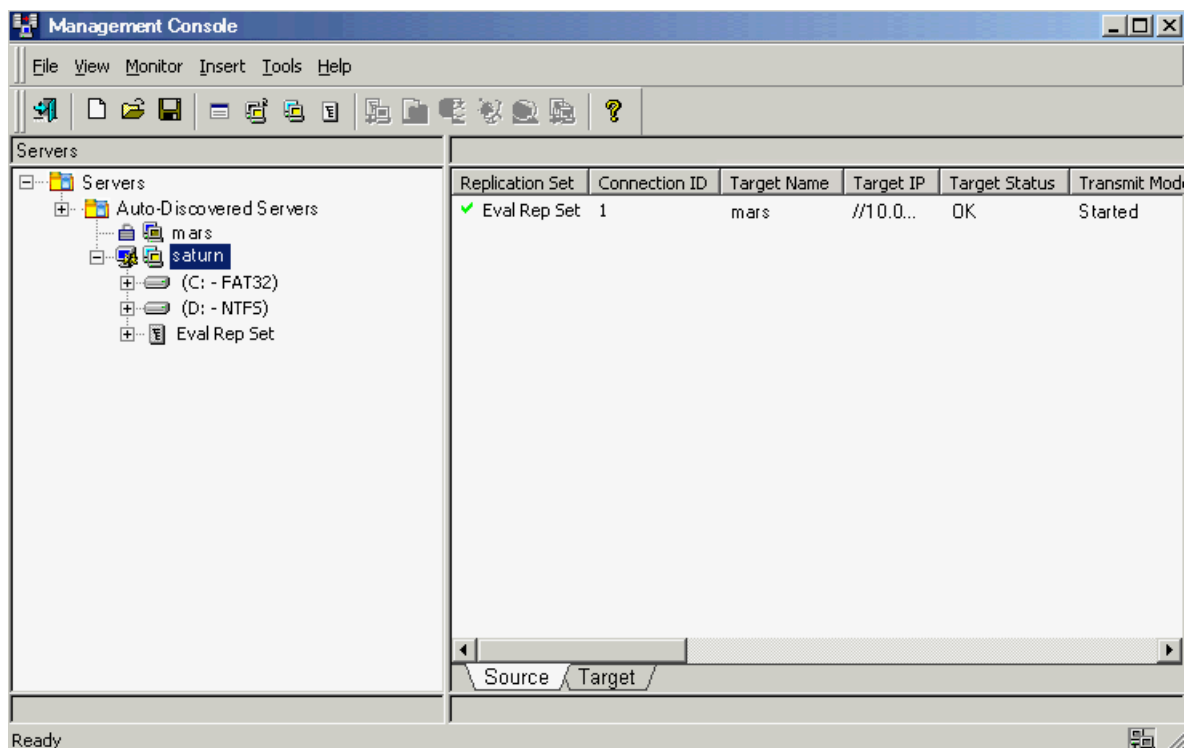
NOTE: These advanced options are explained in full detail in the *Storage Mirroring User's Guide* and are not recommended for this evaluation.

12. You will have the opportunity to configure failover, later in this process, using the Failover Control Center and it is not necessary to make any changes to any of the advanced settings for this evaluation. Therefore, click **Connect** to establish the connection.



NOTE: You will be prompted to save your newly created replication set. Click **Yes** to save it and establish your connection.

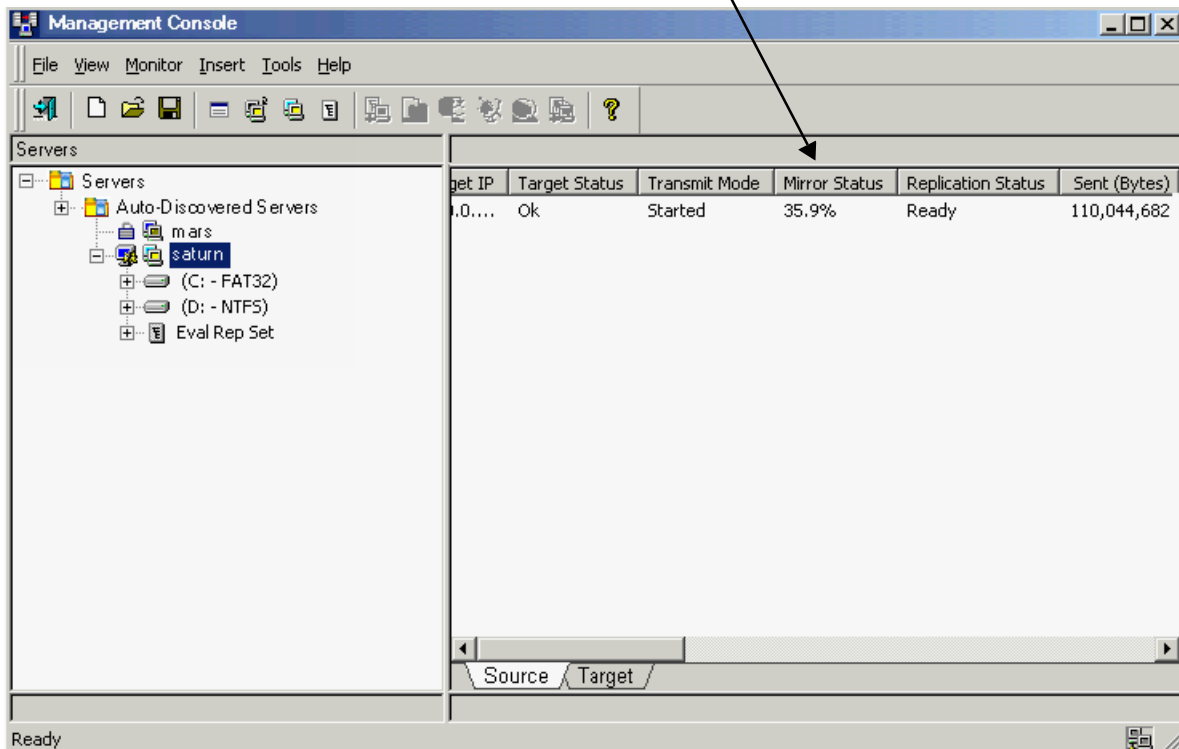
Now that you have created a replication set on the source and connected it to the target, mirroring is initiated. Highlight your source machine on the left pane of the Management Console and you will see your connection in the right pane. For detailed information on understanding the display of the Management Console, see the *Storage Mirroring User's Guide*.



Monitoring the activity and completion of the initial mirror

View your connection in the Management Console by highlight the source machine on the left pane. The connection will appear on the right pane. Use the horizontal scroll bar at the bottom of the right pane to view the various status columns. Pay attention to the **Mirror Status** column which shows the status of the mirroring operation. During the mirroring process, you will see a percentage of the mirror that has been completed. When the **Mirror Status** indicates **Idle**, there is no mirroring activity, meaning your initial mirror has completed. See the *Storage Mirroring User's Guide* for information on the other possible states.

View the status of your mirroring process



To view specific mirroring statistics that may be of interest, use the horizontal scroll bar at the bottom of the right pane of the Management Console window to view the various columns.

- ◆ **Sent (Bytes)**—The total number of mirror and replication bytes that have been sent during this connection.
- ◆ **Sent Mirror (Bytes)**—The total number of mirror bytes only that have been sent during this connection.
- ◆ **Skipped Mirror (Bytes)**—The total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is the same on the source and target machines.
- ◆ **Remaining Mirror (Bytes)**—The total number of mirror bytes only that remain to be sent to the target.

For complete details on all of the statistics, see the *Storage Mirroring User's Guide*.

When your **Mirror Status** is **Idle**, your mirroring is complete. Take a look at your target machine and you will see that the `\testdata` directory exists and is populated with all of the files from your source machine `\testdata` directory.

After your mirror is complete, you are ready to continue with the evaluation.

Changing the test data to cause replication

In order to test replication, you need to change the data contained in your replication set. This includes modifying existing files, creating new files, deleting files, and changing permissions and attributes. However, before changing these files, you will need to do a little bit of preparation work for the verification section.

- 1. From your client machine, browse through the directories of your s: drive.
- 2. On your target, share the new directory that was created by the Storage Mirroring connection. Call the share testshare.
`net share testshare=d:\testdata`
- 3. From your client machine, map t: to testshare on your target.
`net use t: \\source\testshare`
- 4. Browse through your t: drive and verify that it is identical to the s: drive.
- 5. Select four files from your s: drive and record the file information specified in the following table.

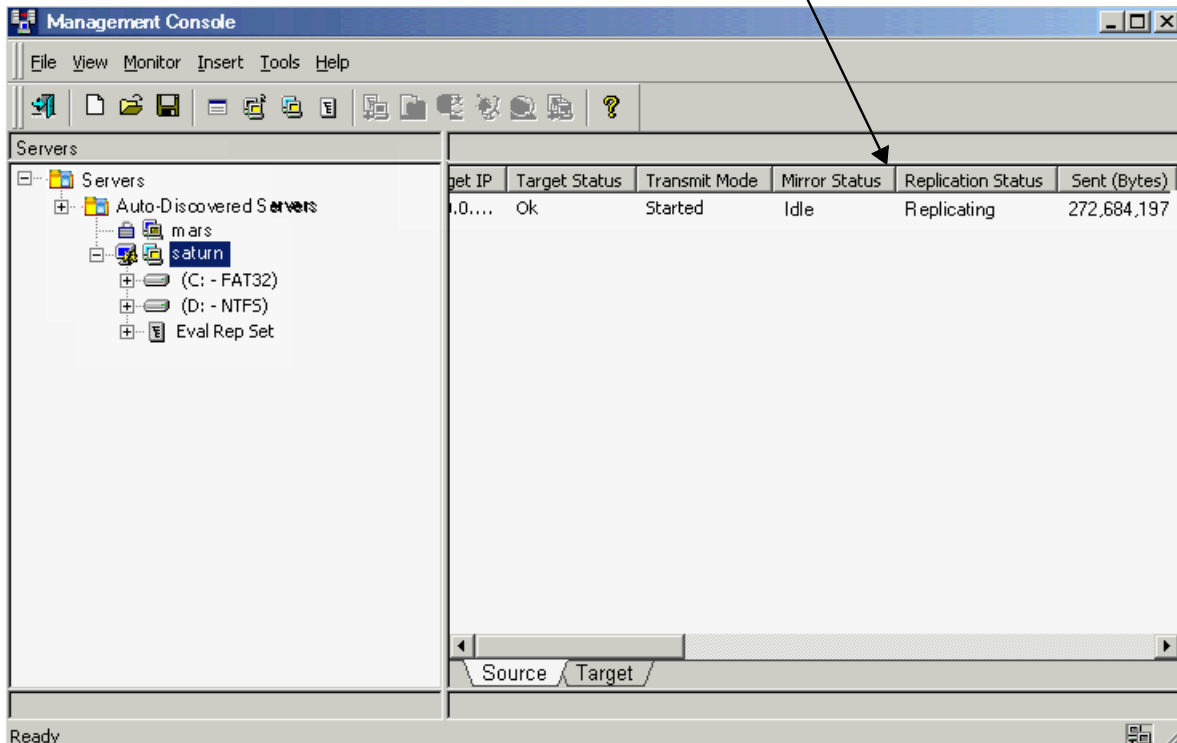
File Name	Attribute to Record
File 1 _____	Date_____ Time _____
File 2_____	Date_____ Time _____
File 3_____	File size _____
File 4_____	File size _____

- 6. On your s: drive, view the contents of one of your files from the table above. Make notes here of the file contents, if necessary.

- 7. On your t: drive, type out the contents of the same file. It matches the file on your s: drive.
- 8. If your Management Console is closed, open it. Tile your Manaagement Console so that you can see it while still having access to your desktop.
- 9. On your s: drive, edit the file that you viewed in step 6. Save your changes.
- 10. Modify the other three files from step 5 so that the date, time, and/or size is updated.
- 11. Watch the Management Console statistics as the file changes cause replication to occur.

12. Just like mirroring, you can monitor the progress of replication through the Management Console. Highlight the source machine on the left pane and the connection information will appear on the right pane. Use the horizontal scroll bar at the bottom of the right pane to view the **Replication Status** column. While Storage Mirroring is actively replicating, the status will be **Replicating**. When there is no replication activity, the status is **Ready**. See the *Storage Mirroring User's Guide* for complete details on all of the replication states.

View the status of your replication process



13. Use the horizontal scroll bars to display additional replication statistics.

- ◆ **Sent (Bytes)**—The total number of mirror and replication bytes that have been sent during this connection
- ◆ **Queued Replication (Bytes)**—The total number of replication bytes that remain in the source queue
- ◆ **Sent Replication (Bytes)**—The total number of replication bytes that have been sent during this connection
- ◆ **Last File Touched**—Identifies the last file that Storage Mirroring transmitted to the target

For complete details on all of the statistics, see the *Storage Mirroring User's Guide*.

NOTE: Many user applications typically save an entire file even though only a portion of the file may have changed. Therefore, the replication statistics will show the entire file being transmitted, not just the changed data. To confirm that replication only transmits the changed segments of files, you must use an application, such as a database application, or a command, such as the echo command, to save only the changed portions of a file.

You may notice your **Replication Status** toggle between **Replicating** and **Ready** as it continues processing the file changes, when your **Replication Status** stays at **Ready**, Storage Mirroring is waiting for additional changes to transmit. After replication is complete, you are ready to continue with the evaluation.

Verifying the changes on the target server

Now that you have modified some of the files, you want to be sure that the file modifications were applied correctly. For this evaluation, three different verification scenarios are presented to demonstrate Storage Mirroring functionality and so that you can be certain that the Storage Mirroring functionality is working as described.

NOTE: Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing, as you are in this evaluation, may have file operations that remain in the cache until additional operations flush them out. This may make Storage Mirroring files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source, the files will be updated on the target. To make sure this does not impact your testing, flush the cache by copying a couple of files from one directory to another and then deleting them.

Visual verification

1. From your client machine, browse your `s:` and `t:` drives. Compare the directory structures and the total number of files.
2. Look again at the four files from step 5 under [Changing the test data to cause replication](#). Verify manually that the changes you made in steps 9 and 10 have been applied to the target copy.

Third-party verification

In this section, you will be using the Windows `comp` command to verify that the changes were applied correctly.

1. Using the command below, you can compare the files on the source and target servers.
`comp s:*.* t:*.*`
2. Compare the results of the `comp` command (using the guidelines below) with the visual and Storage Mirroring verification methods.

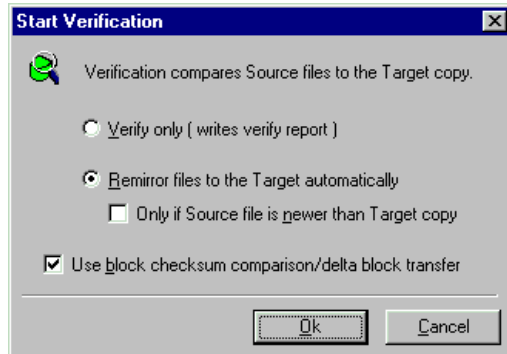
Comp Command Output	Description
Files compare OK	The files are identical
Can't find/open file: <i>filename</i>	The file does not exist on the target
Files are different sizes	The files are different sizes
Compare error at OFFSET 7DCA file1 = 39 file2 = 38	The files are the same size but have different content. The line number which is different is followed by the content, in ASCII, that is different.

```
Comparing S:\file1 and T:\file1...  
Files compare OK  
  
Comparing S:\file2 and T:\file2...  
Files compare OK  
  
Comparing S:\file3 and T:\file3...  
Files compare OK  
  
Comparing S:\file4 and T:\file4...  
Files compare OK
```

Again, you will see that all of the files are reported as the same.

Storage Mirroring verification

1. From your client machine, right-click the connection on the right pane of the Management Console and select **Verify**. You will see two choices on the Start Verification dialog box.

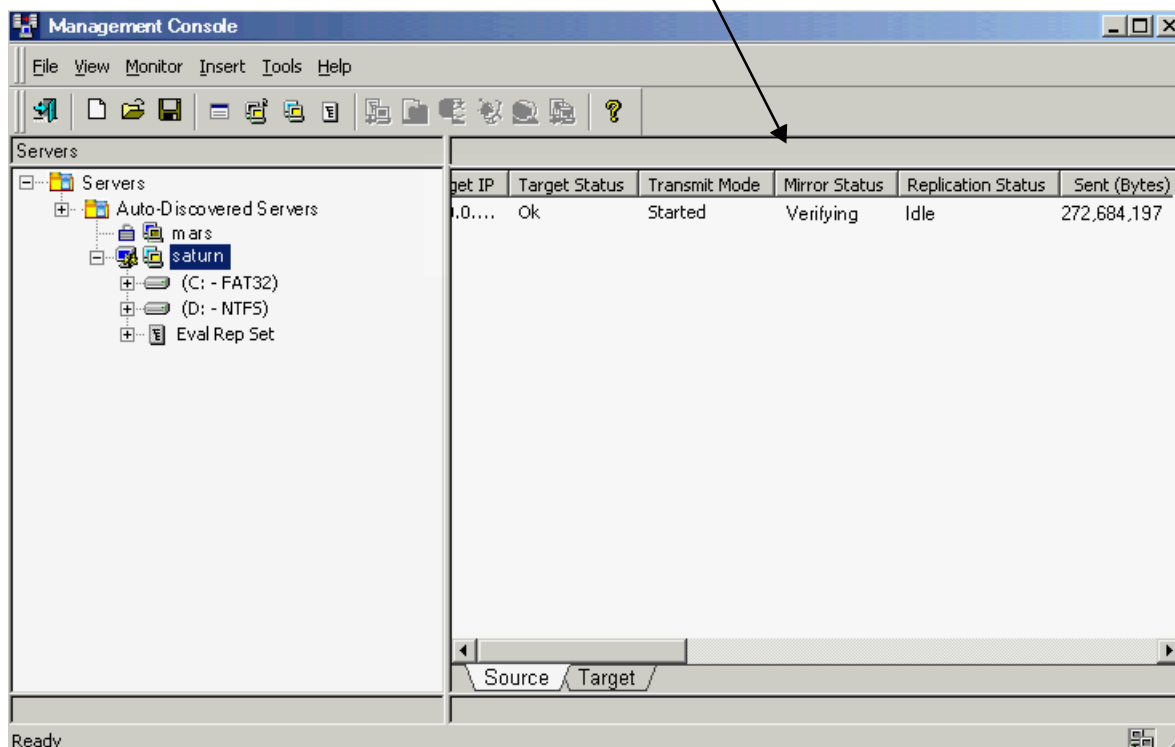


- ◆ **Verify Only**—This option performs the verification process by comparing the date, time and size of each file and generates a verification report identifying the files that are not synchronized.
- ◆ **Remirror Files to the Target automatically**—This options performs the verification process by the comparison method specified, generates a verification rerport, and then remirrors those files from the source to the target that are not synchronized. For complete details on the remirror comparison options, see the *Storage Mirroring User's Guide*.

2. Select **Verify Only** and click **OK**.

Just like when you were monitoring the mirror earlier, you can now monitor the verification process. Notice that the **Mirror Status** column changes to **Verifying** while the verification process occurs. When the verification is complete, Storage Mirroring will have created a log file for you to review.

View the status of your verification process



-
- Wait until your **Mirror Status** has returned to **Idle** and then open the file `DTVerify.log` located in the Storage Mirroring installation directory on your source. You will see that all of the files are reported as the same.

```
Completion Time: 2/7/02 17:04:38:271766
Elapsed Time: 121.689167 seconds
```

```
Total Directories Compared: 28
Total Directories Missing: 0
Total Files Compared: 2567
Total Files Missing: 0
Total Files Different: 0
Total Files Encrypted: 0
Total Bytes Skipped: 43685920
Total Bytes Compared: 43685920
Total Bytes Missing: 0
Total Bytes Different: 0
```

```
----- END OF VERIFICATION -----
```

- Modify one of your files on the target and repeat the verification process, but this time, select **Remirror Files to the Target Automatically**.

NOTE: Since your target file is newer, make sure that **Only if Source file is newer than Target copy** is disabled.

- Using the visual or third-part verification, confirm the synchronization of the file on the target.

NOTE: The remaining sections of this evaluation cover failover, failback, and restore. These processes are used for high availability configurations. These sections are optional, although you are encouraged to step through the procedures.

Configuring failover

When you were looking at the advanced options in the Connection Manager, you noticed a Failover tab allowing you to configure failure monitoring. For this section we are going to use the Failover Control Center so that you can see all of the options and get a feel for how flexible the configuration options are and how easy it is to set up and monitor a source machine.

- From your client machine, open the Failover Control Center (from the desktop, select **Start, Programs, Storage Mirroring, Failover Control Center** or from within the Connection Manager, select **Tools, Failover Control Center**).

NOTE: For detailed information on the Failover Control Center display, see the *Storage Mirroring User's Guide*.

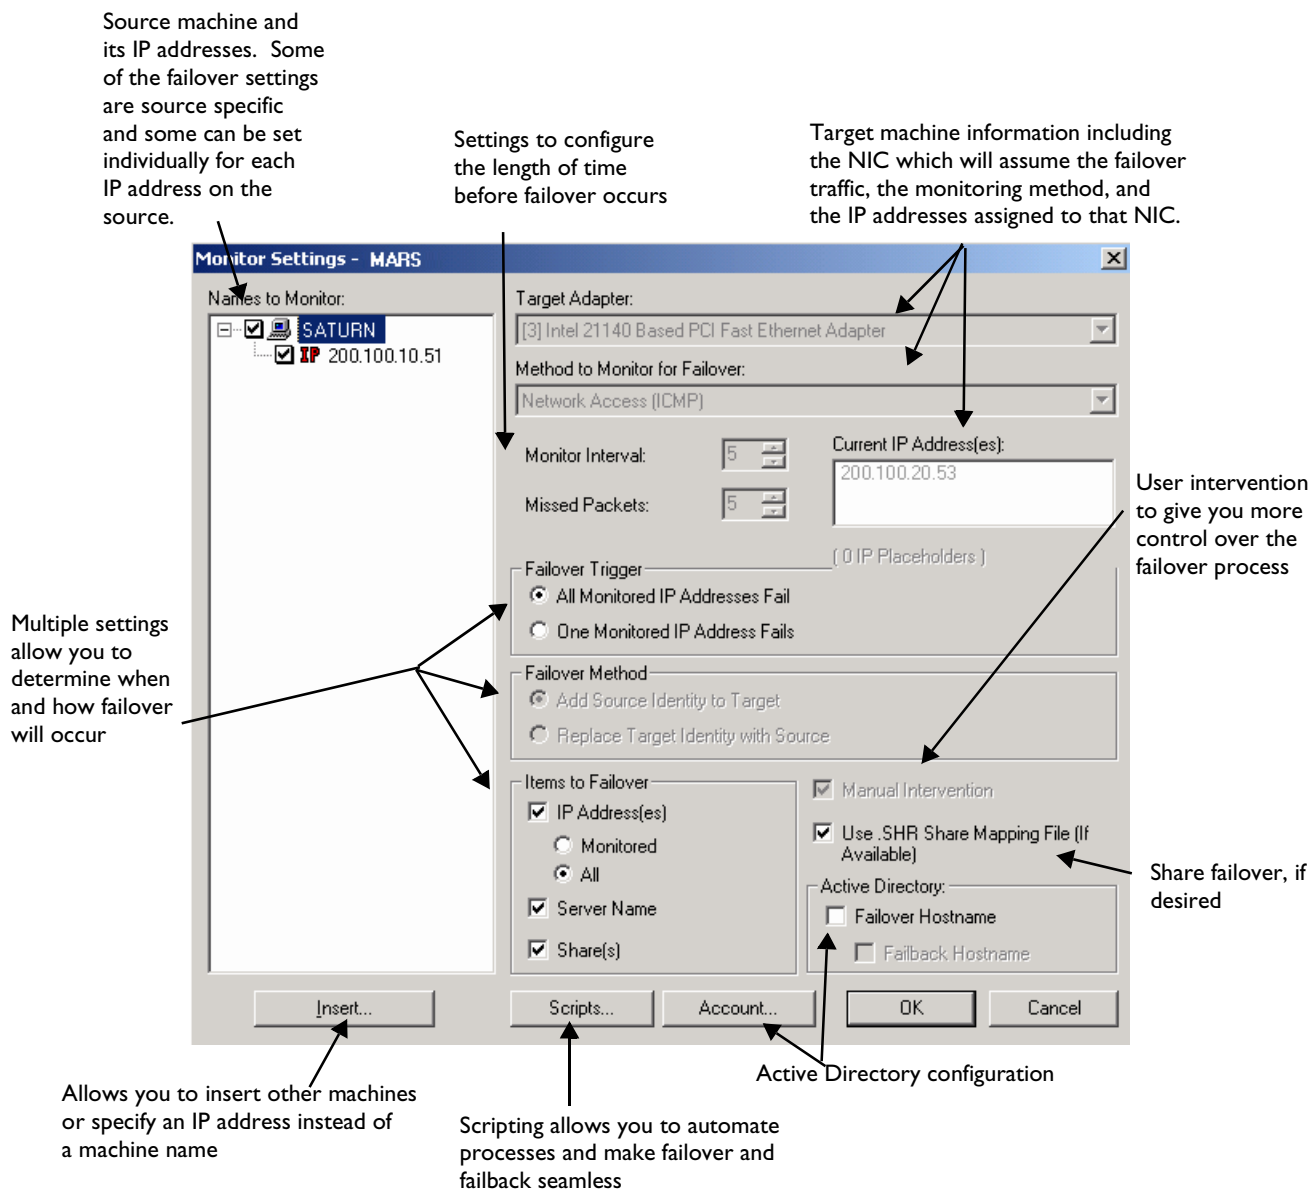
- Select your target from the **Target Machine** list box.

NOTE: You may need to log on to your target machine using your Storage Mirroring administrator user name and password.

If the target you need is not listed, click **Add Target**, and manually enter a name or select the **Browse** button to search for a target machine name.

- Select the source machine to monitor by clicking **Add Monitor**. The Insert Source Machine dialog box appears.
 - Type in your source machine name and click **OK**. The Monitor Settings dialog box appears. This is where you set your monitoring configuration.
 - Select the source to be monitored by marking the check box to the left of the source server name in the **Names to Monitor** tree.
-

At this point, in terms of your evaluation, your failover configuration is complete because you will be using the default settings. But while you are viewing the Monitor Settings dialog box, notice the flexible configuration options available to you. The following diagram points out the key features of the Monitor Settings dialog box. Check the *Storage Mirroring User's Guide* for complete details.

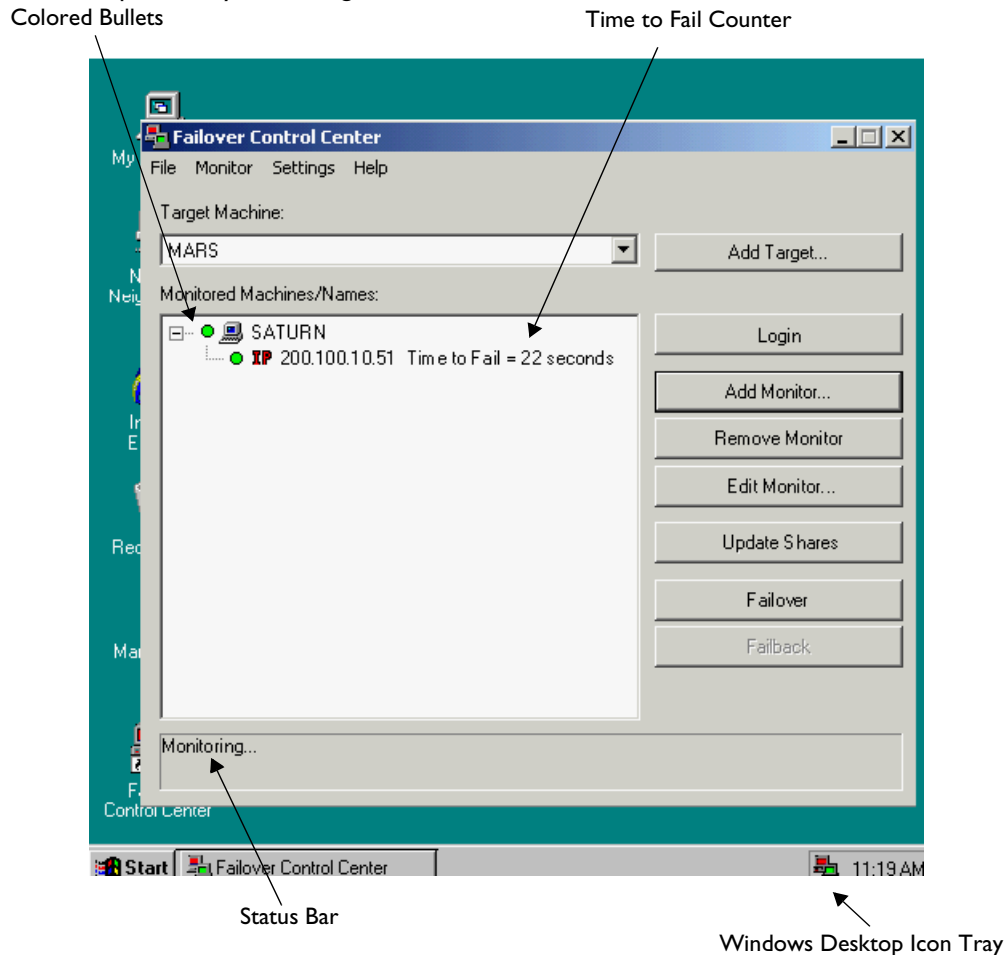


6. After you are done reviewing the different settings, click **OK** to initiate monitoring.

Monitoring failover

Now that failover monitoring is configured and started, you will need to know if and when there is a problem. Since it can be essential to quickly know the status of your machines, Storage Mirroring offers various methods for monitoring the status of failover. When the Failover Control Center is running, you will see four visual indicators:

- ◆ The Failover Control Center Time to Fail counter
- ◆ The Failover Control Center status bar located at the bottom of the window
- ◆ The Failover Control Center colored bullets to the left of each IP address and source machine
- ◆ The Windows desktop icon tray containing a failover icon



NOTE: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

Initiating a failure

If you want to continue with the remaining portions of this evaluation, failback and restore, then you will need to initiate a failure of your source machine. The first few steps of the following procedure will be used as a comparison for after the failover has occurred and then the last steps prepare for the restoration process.

1. Verify that you can view `testshare` on your source.
2. Ping the source's IP address from your client machine.
3. Ping the source's machine name from your client machine.
4. The Failover Control Center does not have to be running in order for failover to occur. For the purpose of this evaluation, make sure that it is running so that you can see each step of the process.
5. Disconnect the network cable on the source. Notice immediately, that the Failover Control Center Time to Fail counter decreases and never resets. You will see the icons change to yellow and eventually to red. Once the icons are red and the Failed Over message is displayed, failover has occurred.

NOTE: The Windows NT Event Viewer on the target provides details on the actual steps that have occurred during failover.

6. Again at your client machine, verify that you can still view `testshare`.

NOTE: Shares may take a few minutes to failover. If they are not immediately available, wait for several minutes and then retry this step.

7. Ping the source's IP address from your client machine.
8. Ping the source's machine name from your client machine.

As you can see, the target has taken on the identity of the source. Application and user requests destined for the source are routed directly to the target. The impact on your end users is minimal.

While your source is failed over to your target, end users were continuing to work without interruption and the data on the target was being updated. After failback, your data on the source is out-of-date. You will want to restore the newer data from the target so that the two machines will again be synchronized.

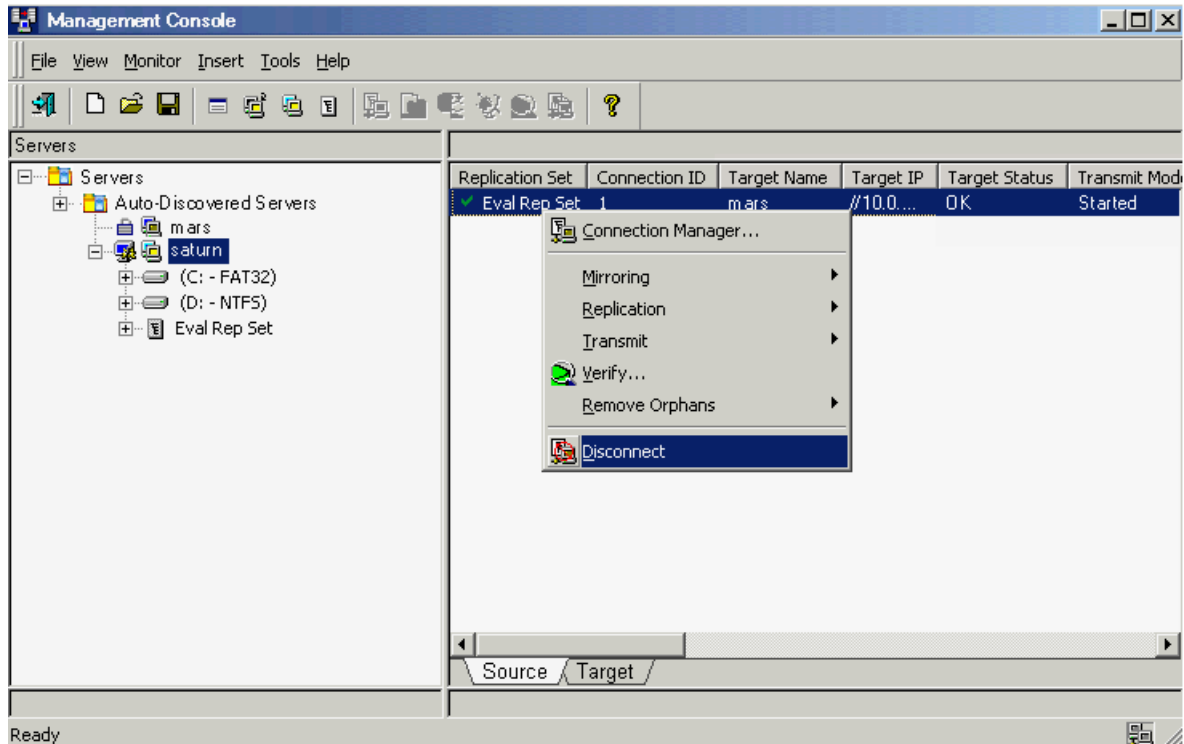
To test the restoration, you need to simulate the changes that the end users would have made on the target while the source was unavailable.

9. Like you did earlier, edit the same file, but this time on the target, that you viewed in step 5 under *Changing the test data to cause replication*. Save your changes.
10. Modify the other three files, on the target, from step 4 so that the date, time, and/or size is updated.

Performing failback

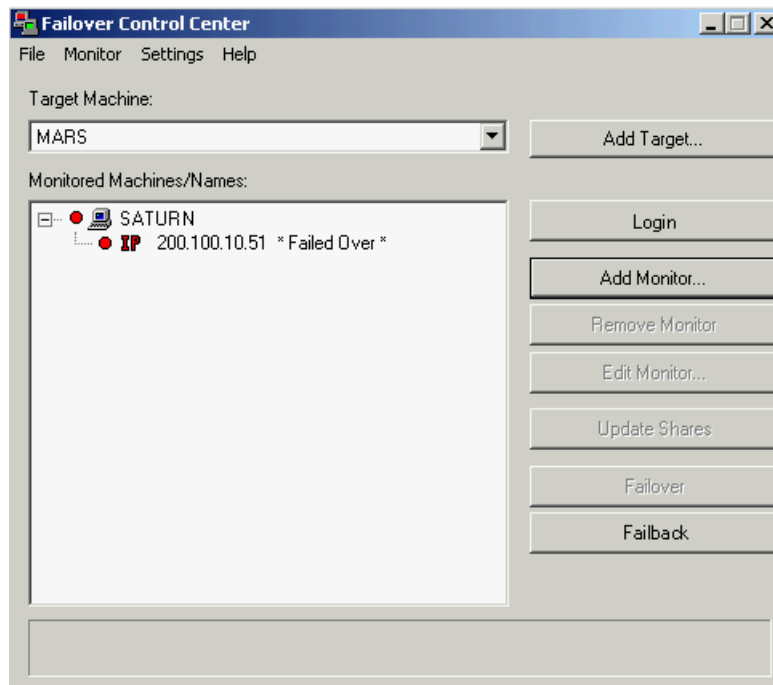
If this was a real-world scenario, you would have the opportunity to fix the problems on the source, without downtime for your end users because the target is standing in for the failed server. When you have resolved the problems on the source, you are ready to perform failback. Failback reverts the target back to its original identity so that the source can be brought online.

1. Before bringing your source machine back online, you will need to disconnect your existing connection. This step is in preparation for the restoration that will be performed in the next section. Open the Storage Mirroring Management Console on the source, login to the source server, right-click the connection, and select **Disconnect**.




2. Now, from the Failover Control Center on your client machine, select the target machine that is standing in for the failed source.

3. Highlight the failed source and click **Failback**.



4. You will be prompted to determine if you want the target to continue monitoring the source. Select **Stop**.



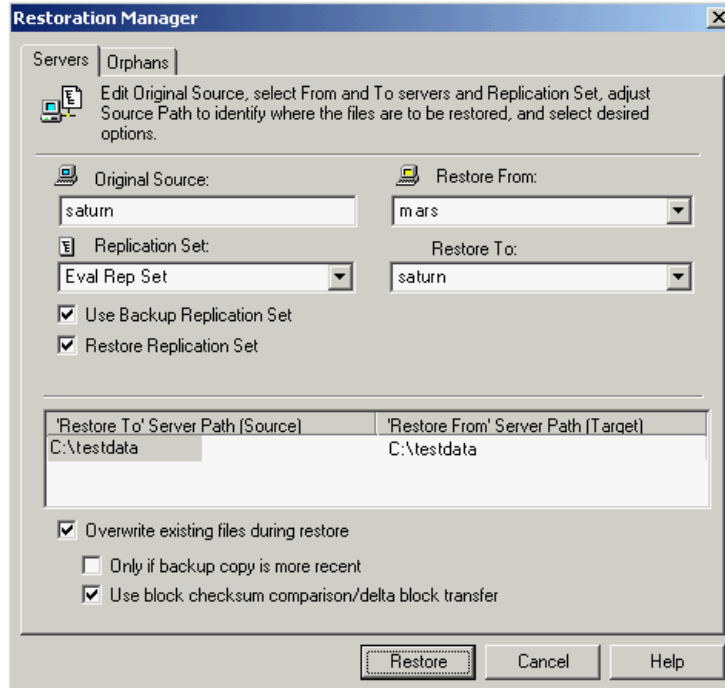
5. Reconnect your network cable on the source machine.

Failback is now complete. If you would like you can repeat the share and ping steps from the previous section to verify that your source and target servers have reverted back to their original identities.

Restoring data

These steps walk you through restoring the newer data from the target so that the source and target are again synchronized. Before beginning the restoration, you will delete one of the files on the source to confirm the restoration process.

1. Delete one of your test files from your source.
2. Open the Management Console on your client machine and double-click your target machine to log in.
3. Select **Tools, Restoration Manager**.

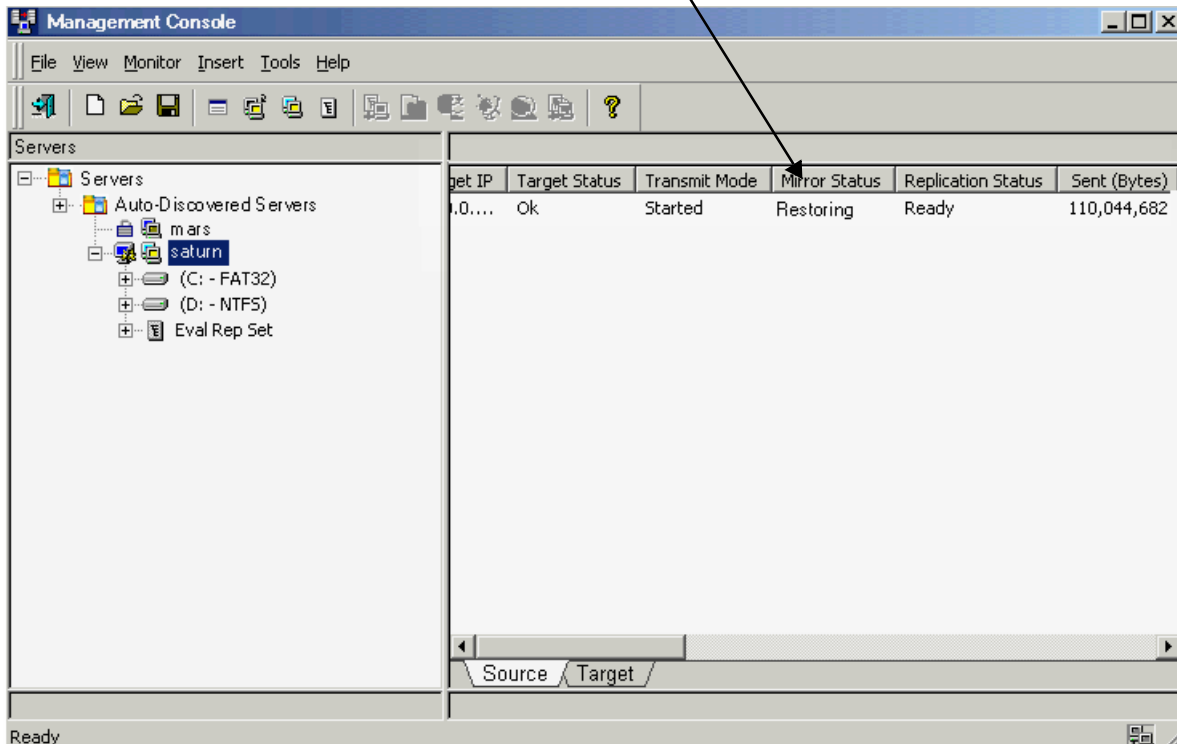


4. Select the **Original Source** machine. This is your source machine where the data originally resided.
5. Select the **Restore From** machine. This is the target machine where the backup data is stored.
6. Select the **Restore To** machine. This is the machine where the backup data will be sent.
7. **Replication Set** contains the replication set information stored on the target machine (the machine in **Restore From**). Select your replication set.
8. The **Source Path** and **Target Path** will automatically be populated when the replication set is selected. Storage Mirroring will allow the data to be placed in a location other than the original, but by default, and for the purposes of this evaluation, restore the data to the same location it came from.
9. For this evaluation, use the default restoration conditionals. For detailed information on the different options, see the *Storage Mirroring User's Guide*.
10. Click **Restore** to begin the restoration.

NOTE: It is important to note that you should not let users access the source or target servers during the restoration process. The restoration process is similar to the original connection process, except that only mirroring is activated. Therefore, the data must not be changing on the target during the restore. Changes made to files on the target during the restore will not be transmitted back to the source. Additionally, files that are in use on the source may prevent the restore from updating those files.

Monitor the verification process as you did the mirror and verification. Notice that the **Mirror Status** column changes to **Restoring** while the restoration process occurs.

View the status of your restoration process



When the restoration is complete, the connection will automatically disconnect. After the restoration connection is disconnected, your source and target machines are synchronized. Confirm that the file you deleted has been re-created by the restoration.

Users can access the source machine and you can, if desired, continue to protect your source by simply reconnecting the existing replication set on your source to the target.

5 ► Conclusion

After you have evaluated the benefits of powerful data protection software from HP, you can explore other ways to enhance and optimize your enterprise solution. Engage HP Sales and Training to help you realize your full potential. HP delivers a comprehensive portfolio of services that help you assess, design, plan, and implement effective data availability and disaster recovery solutions. These solutions help you avoid costly data loss and downtime.

If you have questions on Storage Mirroring, including pricing and product features, call HP Sales or Customer Support.